arXiv:1409.4095v1 [cs.CV] 14 Sep 2014

CAVLECTOMETRY: TOWARDS HOLISTIC RECONSTRUCTION OF LARGE MIRROR OBJECTS

J. BALZER, D. ACEVEDO-FELIZ, S. SOATTO, S. HÖFER, M. HADWIGER, AND J. BEYERER

Abstract. We introduce a method based on the deflectometry principle for the reconstruction of specular objects exhibiting significant size and geometric complexity. A key feature of our approach is the deployment of an *Automatic Virtual Environment* (CAVE) as pattern generator. To unfold the full power of this extraordinary experimental setup, an optical encoding scheme is developed which accounts for the distinctive topology of the CAVE. Furthermore, we devise an algorithm for detecting the object of interest in raw deflectometric images. The segmented foreground is used for single-view reconstruction, the background for estimation of the camera pose, necessary for calibrating the sensor system. Experiments suggest a significant gain of coverage in single measurements compared to previous methods. To facilitate research on specular surface reconstruction, we will make our data set publicly available.

1. Introduction

1.1. **Motivation.** Reconstructing the geometry of an object or scene from a collection of images is an important problem in applications ranging from assisted surgery to additive manufacturing, quality control, augmented reality, just to mention a few. The problem is particularly challenging when objects exhibit complex non-Lambertian reflectance. For instance, specular surfaces cannot be reconstructed without *some* assumption on the surrounding environment. While generic prior assumptions may be sufficient for classification purposes, accurate metrology calls for precise knowledge of the surroundings, that can be thought of as the "illumination". In *deflectometry*, a controllable illuminant (e.g., a liquid-crystal display, LCD) is reflected through the object's surface onto the image plane, see Fig. 2(a). In order to reconstruct a model of the unknown surface from the known illuminant (input) and its measured image output, the input must be sufficiently exciting. Typically, a series of optical code words is displayed, identifying each of the points on the illuminant uniquely. This way, dense correspondence between the domains of input and output signal is established, the so-called *light map*, assigning to a pixel the scene point[1] it sees *via* the specular surface [10].

To make the basic deflectometry procedure viable for industrial applications, an important open problem needs to be solved: Convex objects, no matter how small, act as de-magnifying lenses and thus reflect a significant portion of the environment. For instance, in Fig. 1(a), almost half of the scene becomes visible in a spherical patch of just a few square centimeters. There are essentially two countermeasures: One can either decompose the surface into smaller tractable pieces,

[1]The terms scene, illuminant, and pattern generator are used synonymously throughout this paper.

cf. [3], but depending on the geometry of the object, the number of required views can easily reach the hundreds. Alternatively, many state-of-the-art methods optimize the effective area on the mirror surface by placing the pattern generator in close proximity, cf. Fig. 1(b). Going one step further, Hong et al. [15] arrange five LCDs tightly around the object leaving a small gap in between for a *single* camera to peek through. Both strategies, however, make it difficult if not impossible to account for occlusions, the crux of any reconstruction methodology. In fact, one of the idiosyncrasies of deflectometry is that it may suffer from occlusions not just of the surface itself but also of the *illuminant*, i.e., light rays can be blocked *after* interacting with the surface *and before*. For the purpose of reconstruction, each point on the surface must grant at least one unobstructed view on the illuminant, and ideally, images should be acquired from different camera positions until this condition is completely satisfied. To this end, Weinmann et al. [38] furnish their system with some of the desired flexibility by combining eleven cameras and three LCDs with a turntable to carry the specimen. While their approach yields impressive reconstructions of small-scale objects, first, it still relies on a very large number of measurements, 792 for the example in Fig. 1(c), and second, it excludes large

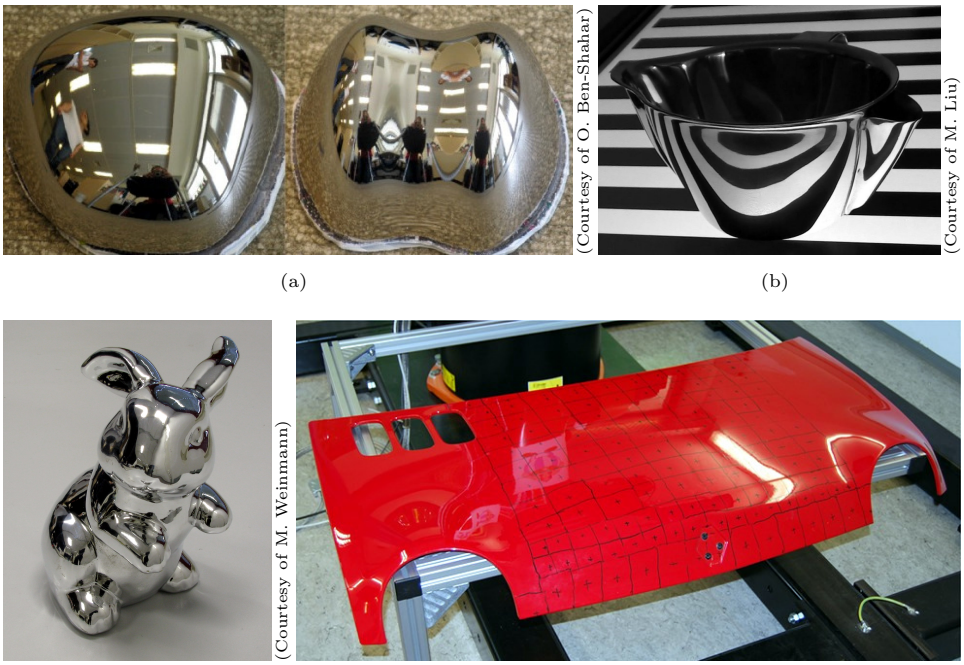

FIGURE 1. Objects of study in some of the current literature on mirror surface reconstruction [1, 20, 38, 3]: (a) Convex surface patches compress the reflected image of the environment. (b) - (c) This effect can be remedied by placing the surface close to the pattern generator and sampling a large number of vantage points. (d) The markings on this engine hood give an impression of how many measurements were required by the method described in [3].

surfaces, e.g. automobile parts, see Fig. 1(d), which are certainly of interest to the practitioner.

1.2. **Contribution and overview.** Here, we propose not only to extend the illuminant to comprise the entire scene, but to do this in a way that supports acquisition of objects more than a meter in length from multiple vantage points. More concretely, we prove the feasibility of performing deflectometry in an *Automatic Virtual Environment* (CAVE), which we dub *cavlectometry*. A CAVE is a room-sized cube consisting of up to six back-projected walls, see Figs. 2(b) and (c). In its default operating mode, users stand inside the CAVE and, while wearing stereo glasses and tracking devices, experience virtual 3-d scenes [8]. "Abusing" it as an illuminant for deflectometry sparks a series of innovations: Rather than treating each of the walls separately, the encoding scheme introduced in Sect. 2.1 – our first contribution – takes into account the sphere-like topology of the CAVE. In Sect. 3, we develop an calibration method which requires no additional data besides that used in later reconstruction, at the heart of which lies an algorithm for separating image regions arising from either Lambertian or specular surfaces. This is our second contribution. Cavlectometry allows us to reconstruct surface patches of unprecedented size from a single vantage point as demonstrated at hand of synthetic and real data sets (Sect. 5.2). Public deflectometry benchmarks are scarce with [3] being the only notable exception. Thus, the final contribution we offer is to distribute all of our data, including raw images and decoded light maps, upon completion of the anonymous review process. Before we explain our method in detail, let us next put it in the context of the existing literature on specular surface reconstruction.

1.3. **Other related work.** Normals of surfaces with hybrid reflection properties can be recovered from purely radiometric considerations [17, 13, 31]. Not before correspondences are available between pixels and the scene points they portray, tools from geometrical optics can be leveraged: identifying a single point light source renders the computation of the light map trivial; then, under known camera motion, an initial point can be expanded into a surface curve by tracking the location where the highlight first appeared [41]. The standard structure-from-motion pipeline can also be enriched as to explicitly take into account specular reflections of a discrete set of scene points [32]. Zheng and Murata [40] progress from isolated point sources towards a one-dimensional concentric illuminant. The accuracy of their approach however remains limited as long as some points on the light source remain indistinguishable. Circular reflection lines suffice nevertheless for special applications such as measuring the cornea in human eyes [12] or surface interrogation [9, 33]. Savarese and Perona [30] study reflections of a two-dimensional checkerboard pattern.

Controllable illuminants were introduced to boost reliability and density of correspondences: The very first active deflectometry setup – due to Sanderson et al. [29] – contained an array of light-emitting diodes (LEDs) which could be switched on sequentially to detect them in the camera image. The acquisition time can be reduced by showing binary codes in parallel [23]. Both papers address the ambiguity by assuming quasi-parallel illumination. A large body of literature on the subject exists in the field of optical metrology starting with [25] that suggests improvements if the LED array is replaced by a commodity computer monitor. Knauer et al. [18] describe the theoretic limits of the method involving phase-shifted sine patterns as

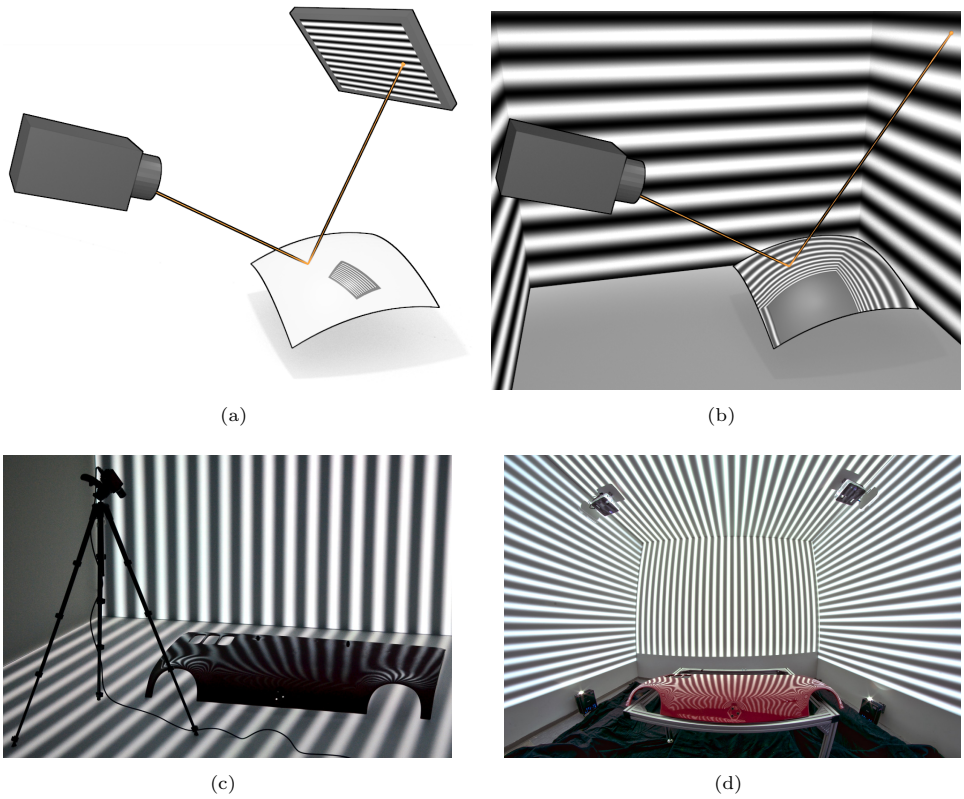

(a) (b)

(c) (d)

FIGURE 2. (a) Current deflectometric measurement setup consisting of a camera and a computer monitor, in between the unknown specular object. (b) Our setup: we place the object in a CAVE which maximizes the encodable scene area *and offers enough space to record images from any desired vantage point.* (c) Cavlectometry recovers the shape of entire car parts like this engine hood from a single view. (d) The principles presented in this paper transfer without modifications to facilities orders of magnitude cheaper to realize than a CAVE.

codes. Light map measurements can be further enhanced by color displays [35]. An active variant of Savarese's method is presented in [28].

Gauge ambiguities can be eliminated by integrability considerations [14, 20] or *specular stereo* [37, 10, 6, 38]. In the special case of infinitesimal displacement of any participant of the imaging chain – camera, surface, or scene – the deflectometric image is perturbed by a *specular flow* field. Blake treats the case of a moving *observer* in a series of papers culminating in [5], which Waldon and Dyer [36] extend by adding photometric effects and interreflections to the model. Like in [41], the trajectories of virtual features in the image plane give rise to reconstructions of one-dimensional subsets of the sought-after surface [24]. They are linked with the caustic curves known in geometric optics [34]. Roth's and Black's [27] method delivers dense two-dimensional reconstructions from specular flow generated by a known camera motion. The authors utilize theoretical results from [7] where, with

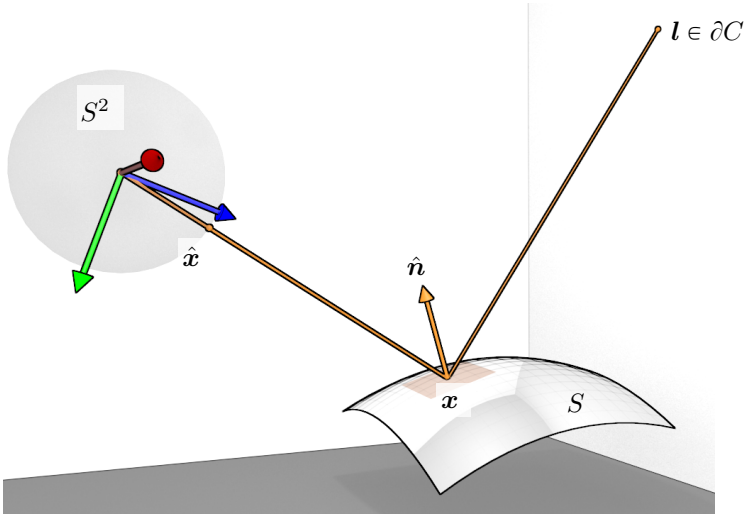

FIGURE 3. Geometry of reflection. The image of the *light map* is the point $\boldsymbol{l}$ that is seen via a point $\boldsymbol{x}$ on the specular surface S under the direction $\hat{\boldsymbol{x}} \in S^2$.

the objective of speeding up ray tracing along piecewise linear paths, it is shown how the points of deflection, i.e., the vertices of the path, behave w.r.t. to changes of the ray emitter. Shape cues may be recovered from the flow field alone, i.e., without image-scene correspondences, when the *scene* is moving relative to object and camera [1]. A closed-form model of specular flow on arbitrarily deforming *surfaces* is derived in [19]. Estimation of specular flow hinges only on the natural scene and can thus handle surfaces of arbitrary size and shape complexity but, as a downside, does not permit metrically precise reconstruction. A comprehensive summary of the state of the art in deflectometry can be found in several recent survey papers [16, 4, 26].

2. Data acquisition

2.1. **Measurement of the light map.** In the following, matrices are in bold, vectors bold italic. Vectors $\boldsymbol{x} \in \mathbb{R}^3$ of unit length are marked by a hat, i.e., $\|\hat{\boldsymbol{x}}\| = 1$. Each pixel in the image plane of a calibrated camera corresponds to a direction $\hat{\boldsymbol{x}}$, under which the associated sight ray leaves the projection center. The relationship between the two is determined by the intrinsic parameters of the camera. Where the ray intersects the surface in $\boldsymbol{x}$, it gets "deflected" according to the geometric law of reflection and travels on until it hits a scene point $\boldsymbol{l}$. For the purposes of geometric modeling, the CAVE is a cube C with boundary ∂C and side length $2h = 3048$ mm. If we neglect multiple reflections and occlusions induced by the camera, the scene point comes to lie on one of the faces of C. Apart from the object under inspection, the scene *is* ∂C and vice-versa. Altogether we obtain a mapping $l : S^2 \to \partial C$, which assigns to a viewing direction $\hat{\boldsymbol{x}}$ the point $\boldsymbol{l}$ seen by the pixel associated with that direction, see Fig. 3. This is precisely the light map defined in Sect. 1.1. Its measurement must be the first step of any deflectometric method.

2.1.1. *Phase shifting.* Although optical phase-shift codes originate from a variety of established applications like optical interferometry, synthetic aperture radar, and magnetic resonance imaging, its features make it the ideal candidate for estimation of the light map: Gradually phase-shifted sine patterns facilitate subpixel accuracy and are relatively insusceptible to disturbances of absolute intensity values, e.g., caused by staining of the specular surface. As in a typical setup (Fig. 2) the camera is focused onto the surface under inspection, limited depth of field exposes the imaged code patterns to blurring. A sine pattern, however, can pass the optical system – which in this case resembles a low-pass filter – unaffected as long its spatial frequency is sufficiently low. In fact, a smoothing of quantization artifacts of the monitor or projector is even advantageous.

Three phase shifts are sufficient to recover a single spatial coordinate uniquely; additional ones may compensate for non-uniformities (e.g., from gamma nonlinearity, etc.). Owing to the fact that phase information is uniquely defined only within the range of one wavelength of the pattern, decoding must be complemented by a so-called *phase unwrapping*. Explaining the codification methodology in detail is beyond the scope of this paper; we must refer the interested reader to [11] among others. We do emphasize here that care must be taken when displaying phase-shifted sine patterns on the walls of the CAVE simultaneously, an issue that never arose in previous systems which were all designed to run on disjoint rectangular screens. The naive approach would be unfolding the cube into the canonical net for cube mapping and shift the patterns along the axes of some local coordinate system. Experimental evidence, however, suggests that treating each face regardless of the others causes problems at the transitions between two maps, see Fig. 4. Here, due to self-illumination, a pattern leaks from one wall into the other. Our key insight is that this is uncritical so long two patterns are symmetric across an edge of ∂C, i.e., if they do not meet orthogonally, at different frequencies or phases. An encoding scheme which fulfills this requirement by moving patterns only parallel to the edges of ∂C is shown in Fig. 5.

2.1.2. *Face encoding.* The inherent structure of this scheme is easily verified: each subsequence encodes one of the three global Cartesian coordinates. Only a single bit of information is missing, prompting a second coding stage. To illustrate the problem, suppose we are looking for the location of a point $\boldsymbol{l}$ on one of the CAVE walls, say $\boldsymbol{l} = (l_x, l_y, l_z)^\top$ where $l_x = h$. Applying the sequences in Fig. 5 will provide the values of l_y and l_z there. The subsequence shifting in x-direction will leave $\boldsymbol{l}$ dark. Putting these facts together, we conclude that $\boldsymbol{l}$ will be on a plane $x = \text{const.}$ but we do not know whether it is the one defined by the equality $x = h$ or $x = -h$. The sign of h is exactly the missing bit. A single image should suffice to set it correctly. The problem of self-illumination, however, persists. In fact, the leakage of a pattern shifting in one of the three coordinate directions into the two orthogonal faces that remain dark meanwhile is even more severe. For this reason, we resort to a complete detection of each of the faces of C. There are altogether six of them prompting the use of a binary code of length 3 bit. An illuminated wall represents a bit value of 1, see Fig. 6. Robustness is further improved with the help of a *differential code*: Thereby, a bit value of 1 is encoded by the *transition* from white to black. Conversely, a bit value of 0 corresponds to an on-switching of the illumination. The differential face coding amounts to in total six images per light map measurement.

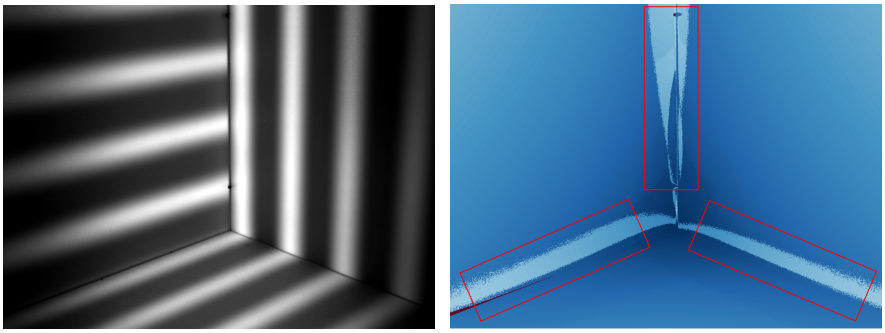

(a) Naive approach

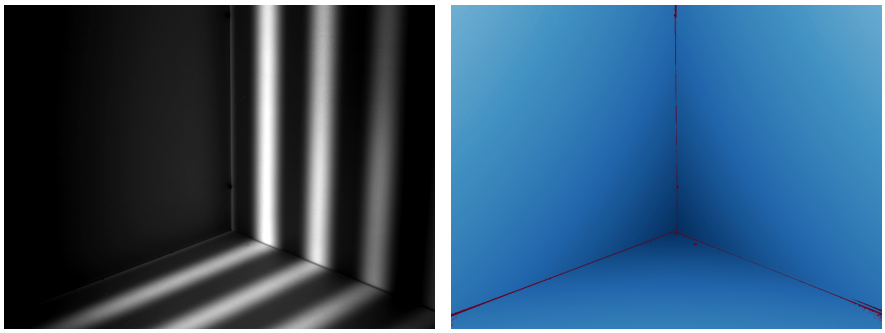

(b) Proposed approach

FIGURE 4. Benefit of a proper encoding strategy: (a) Different walls of the CAVE illuminate each other causing *leakage* of the light map close to the transition between two walls (red boxes). (b) This effect can be mitigated by using only patterns which are symmetric w.r.t. all edges parallel to the shift direction.

3. CALIBRATION

In order to avoid screen occlusions, we advocate recording multiview datasets in a monocular fashion, whereby measurements are taken sequentially while the camera is displaced between two captures. To convert a measured correspondence between $\hat{\boldsymbol{x}}$ and $\boldsymbol{l} = l(\boldsymbol{x})$ into a reasonable normal estimate, both vectors have to be represented in a common reference frame. On one hand, $\hat{\boldsymbol{x}}$ is usually given w.r.t. the frame that is attached to the camera and moves along with it to each new vantage point. On the other hand, $\boldsymbol{l}$ is naturally expressed in terms of the frame at the barycenter of the CAVE (Fig. 5(a) top), which remains static and is thus selected as our world coordinate system. The transformation between the two is an element $g = (\mathbf{R}, \boldsymbol{t})$, $\mathbf{R} \in \mathrm{SO}(3)$, $\boldsymbol{t} \in \mathbb{R}^3$, of the rigid-motion group $\mathrm{SE}(3)$, which – as initially unknown – needs to be determined through calibration.

It turns out that g comes as a byproduct of the very same light map l used for reconstruction later on (Sect. 4). The basic idea of the following procedure is to identify those regions that are imaged *without* intermediate reflections and to localize the camera by means of the appendant scene-image correspondences:

(1) Consider the pullback of the light map $l^* : D \to \partial C$ to the image plane $D \subset \mathbb{R}^2$. It can be calculated knowing the intrinsic parameters of the camera, which we initially estimate by a standard method [39]. The extrinsic calibration process starts with a – possibly very coarse – segmentation of l^*, which builds upon the fact that walls seen directly generate piecewise homographies. One of the distinct properties of homographies is that they maintain cross-ratios, whose computation, however, may be numerically unstable. Instead, we propose a simpler – less discriminative yet effective – inference criterion based on colinearity. Denote by $\boldsymbol{l}_{ij}$ the value of the light map at pixel $\boldsymbol{p} := (u_i, v_j) \in D$ in the discretized image plane. If l^* behaves locally like a homography around $\boldsymbol{p}$, then the points $\boldsymbol{l}_{i,j-1}, \boldsymbol{l}_{i,j}, \boldsymbol{l}_{i,j+1}$ and $\boldsymbol{l}_{i-1,j}, \boldsymbol{l}_{i,j}, \boldsymbol{l}_{i+1,j}$ will be respectively colinear, in other words, $r^x_{ij} = r^y_{ij} = 1$ where

$$r^x_{ij} := \operatorname{rank}(\boldsymbol{l}_{i-1,j}, \boldsymbol{l}_{i,j}, \boldsymbol{l}_{i+1,j})$$

and

$$r^y_{ij} := \operatorname{rank}(\boldsymbol{l}_{i,j-1}, \boldsymbol{l}_{i,j}, \boldsymbol{l}_{i,j+1}).$$

We can define a scoring function $s : D \to \mathbb{N}$ by

$$s(u_i, v_j) := |1 - r^x_{ij} r^y_{ij}|, \tag{1}$$

which vanishes where both rank conditions are satisfied (Fig. 7(a)). A first segmentation of l respectively l^* is obtained by binarizing s with threshold 0.5 (Fig. 7(b)). Following morphological enhancement, this segmentation is refined with the help of a watershed transform [22], see Fig. 7(c).

(2) Observe that there are degenerate cases in which the specular surface itself affords a homography, e.g., if it is locally planar or reflects into a single point (i.e., the light map is constant). We can, however, safely assume that

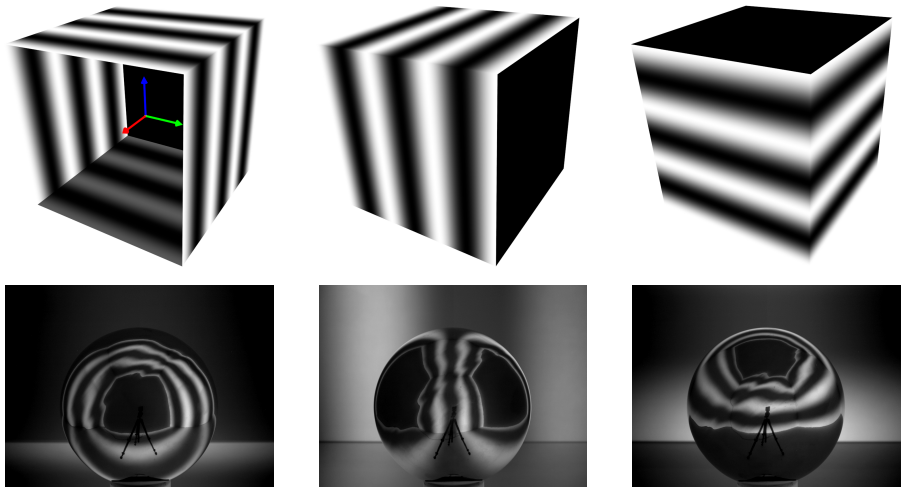

(a) x (b) y (c) z

FIGURE 5. Arrangement of sine patterns shown on the CAVE walls along three spatial directions. During measurement, these are displayed at multiple phases and frequencies.

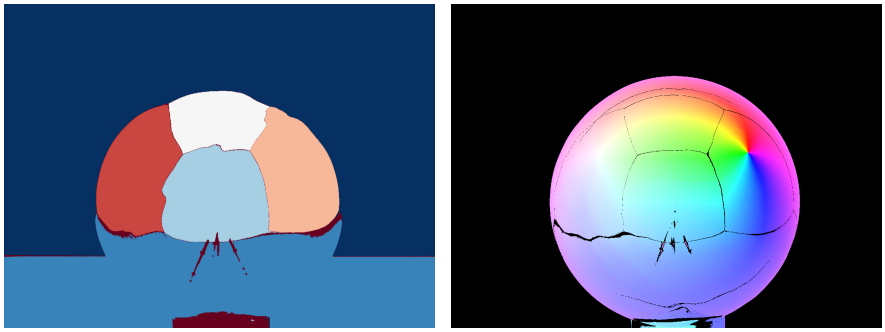

(a) Face ID (b) Full lightmap (view in color)

FIGURE 6. (a) A binary code identifies each separate wall uniquely. (b) Visualization of the full vector-valued light map l. The color map assigns an HSV coordinate to each $\hat{\boldsymbol{l}} \in S^2$ with hue and saturation varying proportionally to azimuth respectively elevation while the value is kept constant at 1.0.

the background occupies a significant amount of the image. Otherwise, we would hardly ever get access to g at all. With this in mind, we shrink the background region determined in the previous step to those pixels that map to the wall which appears most often in it.

(3) The homography between the image plane and the resulting background region, i.e., the CAVE wall that appears in most pixels and is seen directly, is computed by a standard direct linear transformation method. Given the projection matrix, $g \in \mathrm{SE}(3)$ can be extracted from the associated homogeneous matrix [39].

(4) We now revisit the full light map l^* on the image plane. Every point $\boldsymbol{l}(\boldsymbol{p})$ that can be projected directly into the image plane, utilizing the initial estimate of g, must belong to the background. The characteristic function of the foreground then becomes

$$\chi_{\mathrm{b}}(\boldsymbol{p}) := \begin{cases} 1 & \text{if } \|\pi(\mathbf{R}\boldsymbol{l}(\boldsymbol{p}) + \boldsymbol{t}) - \boldsymbol{p}\| < \theta, \\ 0 & \text{otherwise,} \end{cases}$$

where π denotes the canonical pinhole projection. Observe that the sensitivity of χ_b w.r.t. the threshold is quite low. In other words, the choice of θ requires no great care because the backprojection error grows discontinuously across the silhouette. The final result is depicted in Fig. 7(d).

With the final segmentation χ_b, we improve the estimate of the extrinsic camera parameters g by running a robust bundle adjustment on all points in the background [39]. Let us remark that in previous experimental setups, the screen is rarely seen directly during measurement but obviously has to be for extrinsic calibration. To establish an oblique view on the screen, a calibration normal (usually a planar mirror) is necessary, which complicates Step 3 and the final bundle adjustment significantly [18]. Also note that the procedure described in [21] is not applicable here as it requires the pattern generator to be movable.

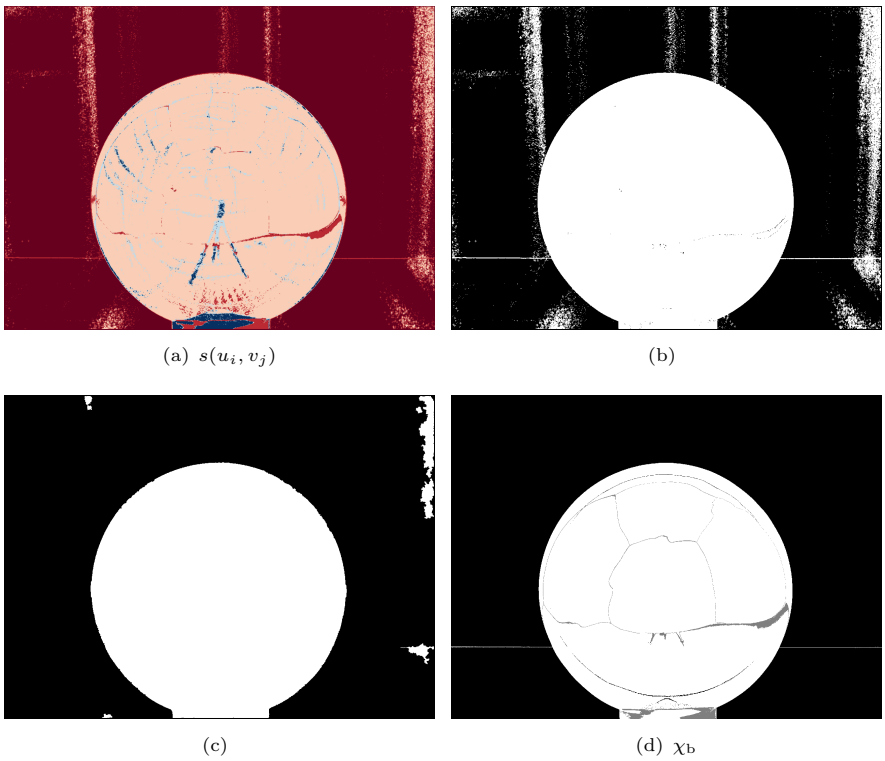

(a) $s(u_i, v_j)$ (b)

(c) (d) $\chi_{\mathfrak{b}}$

FIGURE 7. Segmentation of deflectometric images of strongly-specular objects: (a) Foreground score based on local ranks of the light map. (b) Initial segmentation of (a) by thresholding. (c) Enhancement by watershed transform. (d) The final segmentation is obtained after extrinsic calibration by thresholding the backprojection error.

4. SURFACE RECONSTRUCTION

What sets the deflectometric measurement principle apart from techniques such as stereopsis or laser scanning is that it yields *slope* which must be converted into *depth* by means of *integration*. More precisely, the light map l acquired from a vantage point g is equivalent to a unit vector field of *measured* normals $\hat{\boldsymbol{n}}_{\mathrm{m}}(\boldsymbol{x})$ in space. To see this, fix a point $\boldsymbol{x} \in C$ inside the visible volume, transform it to camera coordinates via $g^{-1}\boldsymbol{x}$, project it to S^2 by normalization, and look up the value of the light map $\boldsymbol{l} = l(\hat{\boldsymbol{x}})$ (Fig. 3). Suppose that $\boldsymbol{x} \in S$, then the surface normal $\hat{\boldsymbol{n}}(\boldsymbol{x})$ must obey the law of reflection, i.e., the correspondence between $\boldsymbol{x}$ and $\boldsymbol{l}$ yields a hypothesis $\hat{\boldsymbol{n}}_{\mathrm{m}}$ about the normal the true surface should have if it indeed contained the point $\boldsymbol{x}$. We say that a regular surface S solves the reconstruction problem if it interpolates the measured normal field in the sense that $\hat{\boldsymbol{n}}(S) = \hat{\boldsymbol{n}}_{\mathrm{m}}|_S$ except on zero-measure subsets of S. Such an S rarely exists due to stochastic disturbances of the data. A maximum-likelihood estimate of the surface is obtain as the solution

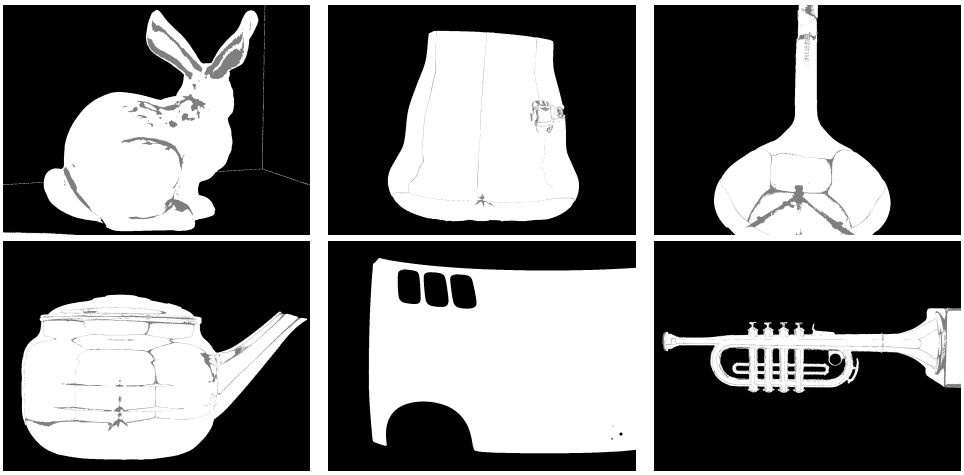

FIGURE 8. Segmentation masks. Pixels with invalid measurements are shown in gray (e.g., on Lambertian parts of the object, the edges of the CAVE, or where the camera occludes the walls).

of the following variational problem:

$$S^* = \arg\min_S \int_C \frac{1}{2} \|\hat{\boldsymbol{n}}(\boldsymbol{x}) - \hat{\boldsymbol{n}}_{\mathrm{m}}(\boldsymbol{x})\|^2 \, \mathrm{d}\boldsymbol{x}. \tag{2}$$

The well-known gauge ambiguity manifests itself in the dependence of $\hat{\boldsymbol{n}}_{\mathrm{m}}$ on $\boldsymbol{x}$: an additional constraint $\boldsymbol{x}_0 \in S$ needs to be imposed to render problem (2) well-posed [4]. Since in this paper, we are mainly concerned with the deflectometric data acquisition process and less with reconstruction, we assume the point $\boldsymbol{x}_0$ to be known for now. The next section will outline a suitable initialization heuristic. For more sophisticated regularization approaches, we refer the interested reader to the related literature, see Sect. 1.3. Also, at this point, we make no attempt to fuse data acquired from a series of vantage points into a common normal field. Our data could be processed with any existing multiview normal field integration scheme, cf. [6, 38].

5. Experimental evaluation

5.1. **Implementation.** We implemented most of the pipeline in C++, except for the steps related to optical coding. All experiments were run on a single 3.4 GHz core of a commodity computer. We solve the shape optimization problem (2) with the algorithm proposed in [2]. Since integration is of broader interest, we will make the source code available. As an initial guess, we use a fronto-parallel plane which is also tangent to the visual hull. Loop subdivision is applied to the evolving mesh between two steps to increase the level of detail and prevent the iteration from falling into local minima of the cost functional[2], see Fig. 10.

[2]although this is less critical given that the method is of second order

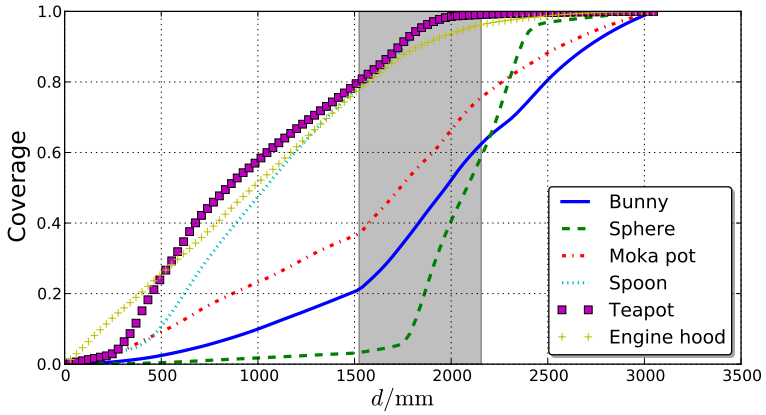

FIGURE 9. Ratio of the area covered by projections of surface points on the image plane (in percentage of pixels) over the area on the screen (parametrized by the distance d to a fixed point) which is seen through these points. Focal length: 16.3 mm.

5.2. **Results.** A selection of segmentation masks obtained with the method developed in Section 3 is shown in Fig. 8. The computation of a single mask constitutes the main bottleneck requiring in average 15 s on a single CPU core. Thereby, the majority of time is consumed by the vast number of singular value decompositions governing the rank estimation in (1). The following experiment investigates the impact of screen enlargement: From the pose previously obtained by calibration, the barycenter of the CAVE wall can be determined which is located *behind* the camera. Fig. 9 shows the number of points in the image of different light maps l closer than a distance d to that barycenter. We limit d to $2h$, as any point further away is likely a part of the background. The shaded area corresponds to the walls next to, above, or below the object, and it can be observed that these may contribute significantly to the number of valid measurements among the pixels of a single deflectometric image. In contrast, the monitor of a classical setup would only occupy a small portion of the back wall.

Our experimental setup is unique in the realm of deflectometry which makes the comparison with other methods challenging. For this reason, we painted some of the

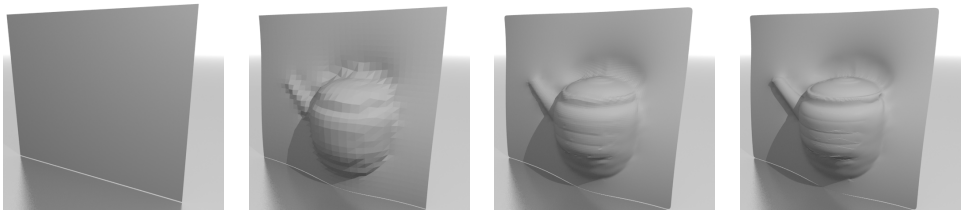

(a) $k = 0$ (b) $k = 1$ (c) $k = 5$ (d) $k = 11$

FIGURE 10. Four iterations of the normal integration method with intermediate uniform refinements of the triangular mesh.

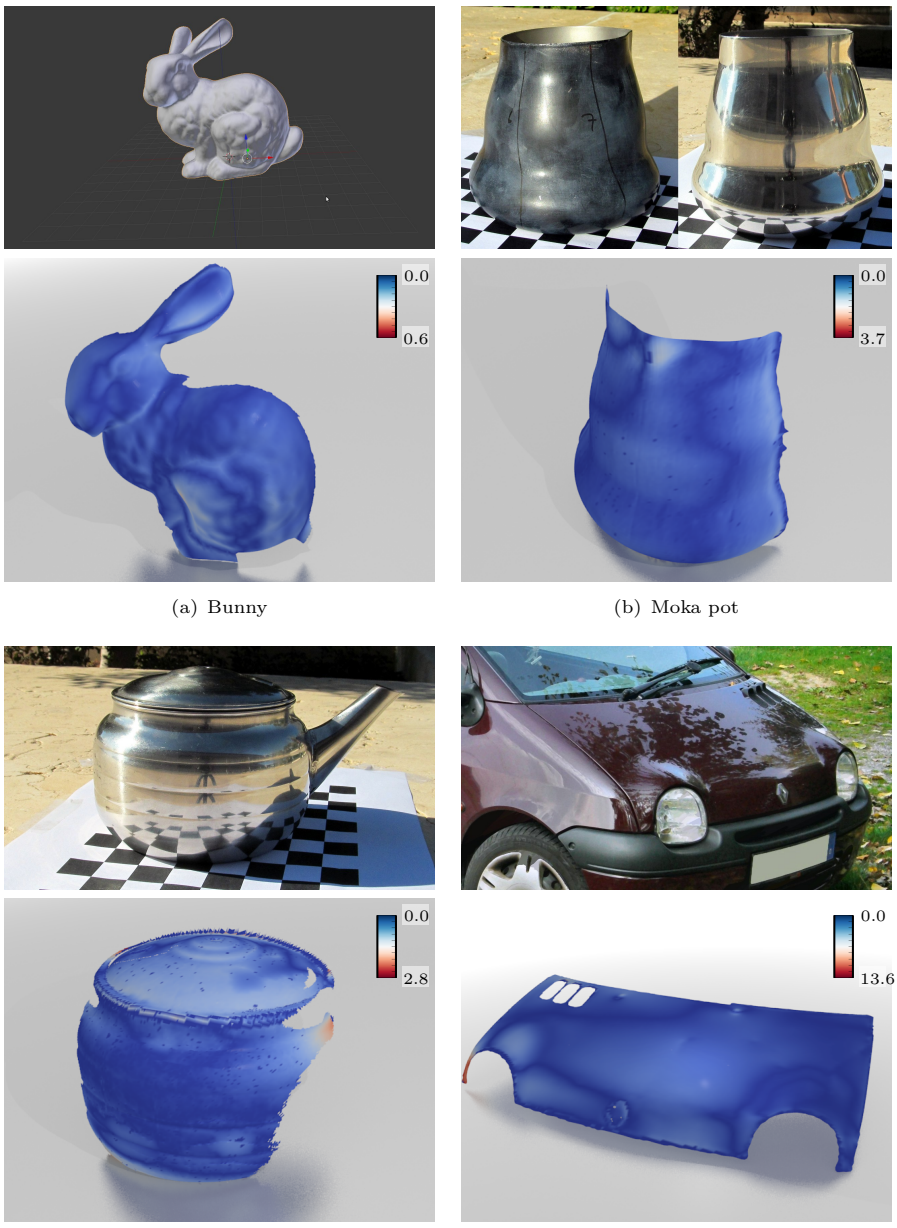

(a) Bunny (b) Moka pot

(c) Teapot (d) Engine hood

FIGURE 11. A picture of the specimen and our reconstruction colored by local deviation from ground truth in mm.

specimens black and coated them with white powder, which is very undesirable in practical applications but unavoidable when using a laser scanner as an alternative source of data (Figs. 11(b) top left). Needless to say that range scans are afflicted with uncertainty of their own. Perfect ground truth is only known after simulating

Object	**No. faces**	**No. views**	**Area** $[m^2]$	**Mean error** $[mm]$	**Max. error** $[mm]$
Bunny	26681	1	0.57	0.002	0.594
Moka pot	25630	1	0.015	0.27	3.69
Teapot	61922	1	0.011	0.39	2.8
Hood	44588	1	0.76	1.12	13.6

TABLE 1. Performance statistics.

data by ray tracing (Fig. 11(a)) or in the form of a CAD model supplied by the manufacturer (Fig. 11(d)). After aligning reconstructed and ground-truth model by the iterative closest point method, the local deviation between the two becomes apparent. Its magnitude is foreshadowed by the vertex colors of the meshes depicted in Fig. 11. In all examples, the error peaks close to the occluding boundary where the actual surface normal is almost perpendicular to the image plane. The mean error is determined by the resolutions of camera and the projectors but also their geometric distance from the surface. The latter is maximized towards better coverage here, which puts the mean error in the range one would expect from a back-of-the-envelope estimate of the backprojected pixel size. Let us remark that our system could be operated in a hierarchical way, where a holistic but coarse reconstruction is examined for abnormalities first, followed by a high-precision close-up scan of conspicuous regions. Aggregate performance statistics are summarized in Tab. 1: Note that *all reconstructions stem from a single capture*. In comparison, Weinmann et al. [38] reported sample sizes ranging from 500 to 750. We were also able to reduce the number of measurements for the same type of engine hood investigated in [3] from 265 to 1.

Finally, we used the recovered meshes and vantage points to render images of the scene they constitute together with a 3-d model of the CAVE. As shown in Fig. 12, the image series leading to the underlying reconstructions could be reproduced with satisfactory accuracy.

6. Conclusion

In this paper, we have argued that deflectometric imaging and reconstruction benefit strongly from illuminants that encompass all of the scene except for the mirror object under inspection. These can be realized e.g. by means of a CAVE. While the cost of building and maintaining such a sensor system may seem limiting, the proposed algorithms transfer – mutatis mutandis – to testbeds far less expensive yet equally effective. The automotive industry, e.g., employs *light tunnels* albeit not for reconstruction. A low-cost alternative in a regular room with opaque white walls and cheap standard projectors is shown in Fig. 2(d). In conclusion, let us remark that the CAVE, beyond its role as deflectometry illuminant, may be utilized for reconstruction in general. To our best knowledge, it has been exclusive to visualization and virtual reality so far, but in our opinion, it opens up possibilities in several other topics of computer vision and optical metrology.

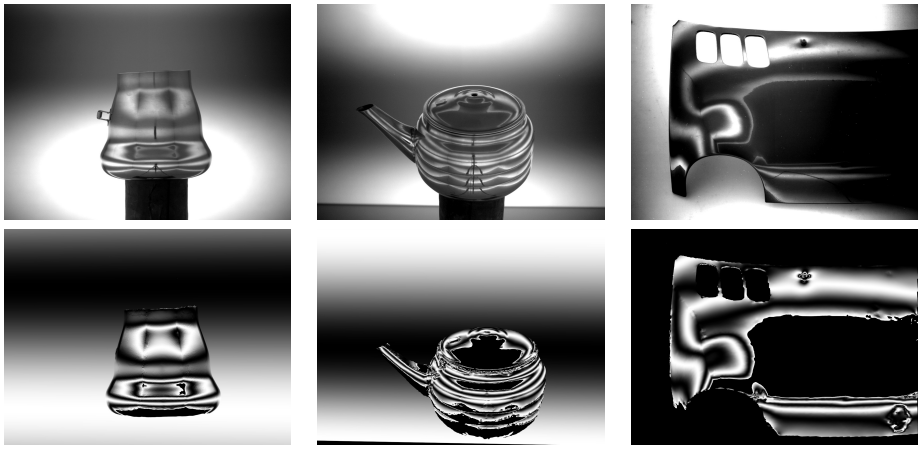

(a) Moka pot (b) Teapot (c) Engine hood

FIGURE 12. Closing the loop: Real images acquired in the CAVE (top) compared to simulations (bottom) based on vantage points and object models obtained by our method.

REFERENCES

[1] Y. Adato, Y. Vasilyev, T. Zickler, and O. Ben-Shahar. Shape from specular flow. *IEEE T. Pattern Anal.*, 32(11):2054–70, 2010.

[2] J. Balzer. A Gauss-Newton Method for the Integration of Spatial Normal Fields in Shape Space. *J. Math. Imaging Vis.*, 441(1):65–79, 2012.

[3] J. Balzer, S. Höfer, and J. Beyerer. Multiview specular stereo reconstruction of large mirror surfaces. *Proc. CVPR IEEE*, 1:2537–2544, 2011.

[4] J. Balzer and S. Werling. Principles of Shape from Specular Reflection. *Measurement*, 43(10):1305–1317, 2010.

[5] A. Blake and H. Bülthoff. Shape from Specularities: Computation and Psychophysics. *Phil. Trans. R. Soc. B*, 331(1260):237–252, 1991.

[6] T. Bonfort, P. Sturm, and P. Gargallo. General Specular Surface Triangulation. *Proc. ACCV*, 3852:872–881, 2006.

[7] M. Chen and J. Arvo. Theory and Application of Specular Path Perturbation. *ACM T. Graphic.*, 19(1):246–278, 2000.

[8] C. Cruz-Neira, D. Sandin, T. DeFanti, R. Kenyon, and J. Hart. The CAVE: audio visual experience automatic virtual environment. *Commun. ACM*, 35(6):64–72, 1992.

[9] Y. Ding and J. Yu. Recovering shape characteristics on near-flat specular surfaces. *Proc. CVPR IEEE*, 1:1–8, 2008.

[10] K. Donner. Stereo Reconstruction of Reflecting Surfaces. *Proc. PRIA*, 14(2):168–172, 2004.

[11] D. C. Ghiglia and M. D. Pritt. *Two-dimensional phase unwrapping: theory, algorithms, and software*. Wiley, 1998.

[12] M. Halstead, B. Barsky, S. Klein, and R. Mandell. Reconstructing Curved Surfaces From Specular Reflection Patterns Using Spline Surface Fitting of Normals. *SIGGRAPH 96*, 1:335–342, 1996.

[13] G. Healy and T. Binford. Local Shape from Specularity. *Comput. Gr. Image Process.*, 42:62–86, 1988.

[14] R. Hicks. Designing a mirror to realize a given projection. *J. Opt. Soc. Am. A*, 22(2):323–330, 2005.

[15] D. Hong, H. Park, and H. Cho. Design of a multi-screen deflectometer for shape measurement of solder joints on a PCB. *Proc. ISIE IEEE*, pages 127–132, 2009.

[16] I. Ihrke, K. Kutulakos, H. Lensch, M. Magnor, and W. Heidrich. Transparent and specular object reconstruction. *Comput. Graph. Forum*, 29(8):2400–2426, 2010.

[17] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE T. Pattern Anal.*, 3(6):661–669, 1981.
[18] M. Knauer, J. Kaminski, and G. Häusler. Phase Measuring Deflectometry: a new approch to measure specular free-form surfaces. *Proc. SPIE*, 5457:366–376, 2004.
[19] J. Lellmann, J. Balzer, A. Rieder, and J. Beyerer. Shape from Specular Reflection and Optical Flow. *Int. J. Comput. Vision*, 80(2):226–241, 2008.
[20] M. Liu, R. Hartley, and M. Salzmann. Mirror Surface Reconstruction from a Single Image. *Proc. CVPR IEEE*, pages 129–136, 2013.
[21] M. Liu, K.-Y. K. Wong, Z. Dai, and Z. Chen. Pose estimation from reflections for specular surface recovery. *Proc. ICCV IEEE*, pages 579–586, 2011.
[22] F. Meyer. Color image segmentation. *Proc. ICIP IEEE*, pages 303–306, 1992.
[23] S. Nayar, A. Sanderson, L. Weiss, and D. Simon. Specular surface inspection using structured highlight and Gaussian images. *IEEE T. Robotic. Autom.*, 6(2):208–218, 1990.
[24] M. Oren and S. Nayar. A Theory of Specular Surface Geometry. *Int. J. Comput. Vision*, 24:105–124, 1996.
[25] D. Perard and J. Beyerer. Three-dimensional measurement of specular free-form surfaces with a structured-lighting reflection technique. *Proc. SPIE*, 3204:74–80, 1997.
[26] I. Reshetouski and I. Ihrke. Mirrors in computer graphics, computer vision and time-of-flight imaging. *Lect. Notes Comput. Sc.*, 8200:77–104, 2013.
[27] S. Roth and M. Black. Specular Flow and the Recovery of Surface Structure. *Proc. CVPR IEEE*, 2:1869–1876, 2006.
[28] S. Rozenfeld, I. Shimshoni, and M. Lindenbaum. Dense mirroring surface recovery from 1D homographies and sparse correspondences. *IEEE T. Pattern Anal.*, 33(2):325–37, 2011.
[29] A. Sanderson, L. Weiss, and S. Nayar. Structured Highlight Inspection of Specular Surfaces. *IEEE T. Pattern Anal.*, 10(1):44–55, 1988.
[30] S. Savarese, M. Chen, and P. Perona. Local Shape from Mirror Reflections. *Int. J. Comput. Vision*, 64(1):31–67, 2005.
[31] H. Schultz. Retrieving shape information from multiple images of a specular surface. *IEEE T. Pattern Anal.*, 16(2):195–201, 1994.
[32] J. Solem, H. Anaes, and A. Heyden. A Variational Analysis of Shape from Specularities using Sparse Data. *Proc. 3DPVT IEEE*, 1:26–33, 2004.
[33] P. Sturm. Recovering specular surfaces using curved line images. *Proc. CVPR IEEE*, 1:2326–2333, 2009.
[34] R. Swaminathan, S. Kang, R. Szeliski, A. Criminisi, and S. Nayar. On the Motion and Appearance of Specularities in Image Sequences. *Proc. ECCV*, 1:508–523, 2002.
[35] M. Tarini, H. Lensch, M. Goesele, and H.-P. Seidel. 3D Acquisition of Mirroring Objects. *Graph. Models*, 67(4):233–259, 2005.
[36] S. Waldon and C. Dyer. Dynamic shading, motion parallax and qualitative shape. *Proc. WQV IEEE*, pages 61–70, 1993.
[37] Z. Wang and S. Inokuchi. Determining Shape of Specular Surfaces. *Proc. SCIA IAPR*, 1:1187–1194, 1993.
[38] M. Weinmann, A. Osep, R. Ruiters, and R. Klein. Multi-view Normal Field Integration for 3D Reconstruction of Mirroring Objects. *Proc. ICCV IEEE*, pages 2504–2511, 2013.
[39] Z. Zhang. A flexible new technique for camera calibration. *IEEE T. Pattern Anal.*, 22(11):1330–1334, 2000.
[40] J. Zheng and A. Murrata. Acquiring a Complete 3D Model from Specular Motion under the Illumination of Circular-Shaped Light Sources. *IEEE T. Pattern Anal.*, 22(8):913–920, 2000.
[41] A. Zissermann, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image. Vision. Comput.*, 7:38–42, 1989.